

Sub
A1

~~MULTIPLEXOR TOOL FOR DISTRIBUTED SYSTEMS MANAGEMENT~~

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to computer network systems and more particularly to a module which permits simultaneous communication between a client program and multiple server programs without requiring different protocols or different versions of client or server programs.

Description of the Related Art

To enable communication between program C (client program) and program S (server program), possibly running on different execution domains (processors), a protocol P is used. Using the protocol P, the client program C may send a request Rq to the server program S, and the server program S is expected to return a response Rp. This abstraction describes a large number of widely available protocols (e.g., HTTP (Hyper-Text Transfer Protocol), RPC (Remote Procedure Calls), RMI (Remote Method Invocation), etc.). This situation, in which a client program C 10 uses a protocol P 11 to communicate with a server

program S 12, is shown in Figure 1. The protocol P, in this situation, is a one-to-one, synchronous protocol.

However, it is not convenient (or even practical) if the same task is to be performed on two, ten or perhaps even hundreds of instances of the server program S (which conceivably could be running on an equal number of network-connected processors) over a one-to-one protocol P. Therefore, a broadcast or multicast protocol P' is used such a situation. Multicast and broadcast protocols distribute one message among a large number of nodes in a network and require a mechanism for handling multiple responses. In general, this is accomplished by modifying the client program C, server program S, or both, such that they work specifically with multicast protocol P'. Thus the conventional solution to permit such one-to-many situations requires modifying the client program C (and most likely the server program S) to use this protocol P' and handle the multiple responses it may receive.

SUMMARY OF THE INVENTION

It is, therefore, an object of the present invention to provide a structure and method for processing a one-to-one request from a client program to multiple instances of a server program over a protocol. The method includes transferring the request from the client program to a multiplexor, generating a plurality of instances of the request using the multiplexor (wherein each of the instances of

the request corresponds to a different instance of the server program), transferring the instances of the request from the multiplexor to the instances of the server program, transferring a plurality of responses from the instances of the server program to the multiplexor, converting the responses to a uniform response; and
5 transferring the uniform response to the client program.

The method specifies target instances of the server program to form a fan out target list, to which the request will be transferred. The converting comprises selecting an operation to combine the responses, such as listing the responses, aggregating the responses, adding the responses, preparing a subset of the
10 responses, identifying a maximum of the responses, identifying a minimum of the responses, and averaging the responses.

The multiplexor automatically creates the instances of the request. The client program, the instances of the server program, and the protocol are unaffected by the method. The unified response has an instance corresponding to
15 the client program.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

20 Figure 1 is a schematic diagram of a one-to-one protocol;

Figure 2A is a schematic diagram of the inventive multiplexor;

Figure 2B is a schematic diagram illustrating the ability of the invention to operate with multiple instances of client and server programs;

Figure 3 is a schematic diagram of the inventive multiplexor;

Figure 4 is a web browser page illustrating the invention;

5 Figure 5 is a web browser page illustrating the invention;

Figure 6 is a web browser page illustrating the invention; and

Figure 7 is a hardware embodiment of the invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS OF THE INVENTION

10 The invention includes a software "multiplexor" that acts as an intermediary in the communications between one client program and multiple instances of a server program. The concept of multiplexing signals from one source to different destinations has been in use in the field of electronics and the invention derives its name from this electronic component by analogy.

15 As discussed above, with the conventional systems the client program, the protocol and/or the server programs are modified to allow a request from the client program to be issued to multiple instances of the server program. This process is very time consuming and expensive. The invention is fundamentally different because the invention only modifies the request (and the response to the
20 request) to accommodate the different instances of the server program. By

modifying only the request and the response, the invention avoids the need to modify the server program, client program or protocol. Further, because only the request and the response are modified, the process can be performed automatically using the inventive multiplexor, which makes the request and response transparent (e.g., appear as a one-to-one communication) to the server and client programs.

A single client program C often needs to send the same request R_q to several instances of a server program S (S_i) and process the responses R_{pi} obtained from each instance of the program S_i . Each server S_i is assumed to be executing the same program. However, the different instances of the server program S_i have different data and/or states.

In computer science the term "program" refers to the implementation of an algorithm in a given programming language. A "process" refers to a program running on a computer. However, the two terms are often interchanged. For consistency and ease of understanding, the term "program" is used throughout the application. However, one ordinarily skilled in the art would understand whether a program or process was being performed, based on the context of the issue being discussed.

Conventionally, the client program C would have to be customized to accommodate each different instance of the server program S_i . The inventive multiplexor M eliminates the need to customize the client program C and automates the communication without modification of the programs C or S_i , or the protocol P.

The invention comprises a multiplexor M, with extensions to handle specific protocols. Further, the invention permits context-free operation under certain assumptions, allowing a single multiplexor to handle communications between many pairs of C-S, C'-S', C''-S'' that share the same protocol P, without
5 any additional programming. As shown in Figure 2B, the inventive multiplexor 22 can operate with different instances of the client program C, C' and different instances of the server program S0, S1, S2, S0', S1', etc., as well as operate with multiple instances of both programs simultaneously.

The invention can also be used in the one-to-one situations in which a
10 single client program C uses a protocol P to communicate with a single server program S, as shown in Figure 1.

The one-to-many case and the role of the inventive multiplexor is illustrated in Figure 2A. In Figure 2A, the client program C 20 can communicate with n instances of the server program S_1, S_i, S_n 24 through the use of the
15 multiplexor M 22 using protocol P 21. The multiplexor's client MC 25 supplies two additional inputs to the multiplexor, the fan-out and the operation op. The fan-out specifies the target servers from among all possible instances of the server program S_i (from a possibly infinite number of instances), while the operation op determines how the results R_{pi} returned by each of the instances of the server
20 program S_i will be combined and made available to the client program C. The multiplexor uses this information to generate multiple instances of the request from the single original request sent by the client program. The servers will send

back a response to the client program, either indicating an error condition or successful execution, possibly returning some data. These responses are then modified and combined by the multiplexor to correspond to the instance of the client program (e.g., to the same format, version, data structure, etc. of the original request) so that the client program believes it is talking to a single server program in a one-to-one communication environment.

Therefore, the multiplexor's client MC determines which instances of the server program S will receive the request from the client C. Additionally, the multiplexor's client MC determines how the responses from the different instances of the server program S will be combined. For example, the responses could be averaged, a minimum or maximum response could be found, all different responses could be listed, a subset of the responses could be prepared (based on some criteria supplied by the multiplexor's client), the responses could be aggregated (again based on a special aggregation formula supplied by the multiplexor's client), or the responses could be added together (e.g., totaled), depending upon the operation op.

Figure 3 illustrates the inventive multiplexor in greater detail. More specifically, the multiplexor M 22 includes a buffer unit 30 which stores requests from the client 10 until the various instances of the server program 24 become available (e.g., come on-line) or until some other significant event takes place, such as a more convenient time of day when system usage is lower. In addition, the buffer 30 can store responses until all responses are received from the various

instances of the server program 24. A syntax converter 32 similarly converts the request into a syntactically correct request for each given server program 24. A response combiner 34 acts under the instructions of the multiplexor's client 25 to combine the responses from the various instances of the server program 24. The response combiner 34 may utilize the syntax converter 32 to modify the responses received from the instances of the server program 24 back into a data format and syntax which is acceptable to the client program 20.

The multiplexor 22 accommodates multiple instances of a protocol by recognizing that there are some commonalities, i.e., some basic set of characteristics that are common to all protocols, such as the operations on results returned by servers, target selection and connection management, that need only be implemented once. Then, to implement different syntaxes and requirements of each protocol, the basic set of features is extended with protocol-specific functions (e.g., addressing, flow and error control, etc.). Because standard protocols are widely used, the extensions to the multiplexor to support one of these protocols can be leveraged across a large number of applications, i.e., for many pairs of the client-server programs C-S.

Thus, the invention allows the designers and developers to concentrate on the simpler, one-to-one protocol interaction and benefit from the generic facilities of the multiplexor.

All relevant protocols provide some mechanism for directing messages from a source to a destination, usually referred to as an "addressing scheme." This

is reflected in the syntax of protocol messages, however the way in which the addressing scheme is implemented will vary from one protocol to another. Since the inventive multiplexor (which acts as an intermediary) redirects requests, the inventive multiplexors must modify the addressing information contained in the request originated by the clients and in the response returned by the server.

Additional syntax conversion may be performed by the invention as required by specific protocols.

The following outlines the typical sequence of events that take place when the client communicates with several instances S_i of the server program using the inventive multiplexor M .

First, the fan-out and operation op are selected by the multiplexor's client MC . For instance, the inventive multiplexor M may be integrated with a systems management program SM that manages several remote servers.

Also, the systems management program SM could provide a graphical user interface to the client program C (e.g., to the user) that depicts each server using an icon and/or groups of servers by some criteria such as location (as shown in Figures 4-6, discussed below). The user can select target servers to be added to the fan out list by clicking on their representation. The fan-out list is passed to the multiplexor's client MC and then to the multiplexor as a parameter. The operation could be similarly selected by the user through the systems management program using the graphical user interface, as also discussed below. While this example provides that the client program C provides inputs to the system

management program 26, the system management program can be controlled independently through a separate user interface (or by a separate network), as would be known by one ordinarily skilled in the art given this disclosure.

Servers with dynamic internet protocol (IP) addresses and sporadic
5 connectivity may communicate with the system manager 26 to inform the system manager 26 when they become available (before they disconnect), as well as provide any dynamic information. Although the system manager 26 is not strictly part of the multiplexor (i.e. they can exist independently), both modules communicate in order for the multiplexor 22 to be able to send the actions to the
10 targets 24. For example, a configurable agent can be used on the target servers. The agent will notify the connection manager when a dialed connection is established, and can also be manually triggered.

When results are available, they can be displayed using any graphical user interface (e.g., Web browser). The user interface is essentially the same as for the
15 one-to-one case, but may represent a combination of the data coming from multiple servers.

Second, the client program C sends a request Rq to the multiplexor M over the protocol P 21. In order to do this, the multiplexor M must accept the same requests that an instance of the server program S would. This, in most cases, is
20 implementation-specific. The syntax converter 32 modifies the request from the client program 20 in different ways so that each of the different instances of the server programs 24 view the request as a one-to-one communication from an

individual client program C. Therefore, with the invention, it appears to each different instance of the server program 24 that the request is coming directly from the client server 20 and not from the multiplexor 22. Thus, the invention permits the client program 20 to send a request to multiple instances of a server program 24 without modifying the client program 20, the server programs 24, or the protocols 21, 23.

The sophistication of the syntax converter 32 and the characteristics of the specific protocol P affect the degree to which the multiplexor 22 can be automatically initiated to handle arbitrary instances of the server programs 24.

For example, when the invention is used with the RMI and HTTP protocols very little or no customization is required to handle each specific server program because the protocols provide enough information for the inventive syntax converter 32 to automatically extract the necessary elements (e.g., addressing as described above) to re-route the requests and combine the responses.

When unusual protocols and/or instances of the server program are encountered (and even if the source code for a server program is not available) the inventive multiplexor can still create the modified requests by observing the format of messages accepted by the target servers. This requires no semantic interpretation of messages, just syntactic compatibility. This makes it possible to develop context-free multiplexors.

Next, from the fan-out list, the multiplexor M identifies all targets instances of the server program S_i through S_n . Then, for each server program S_i ,

the multiplexor M sends the same request Rq that it received from the client program C, using protocol P. Once again, it appears to the instances of the server program Si receiving the request Rq that the request comes from an instance of the client program C.

5 The Protocol P is assumed to be shared by all nodes of the network, and is assumed to handle only one-to-one communications. As described above, from the point of view of the server programs, the multiplexor creates "virtual" instances of the client program C by modifying the request. Conceptually, an infinite number of simultaneous connections could be supported.

10 Each server program Si processes the request and returns a response Rpi to the multiplexor M, which stores the response Rpi in the buffer 30. Once all responses Rp1...Rpn have been received, the multiplexor 22 uses the processor 35 to apply the operation op to the response list. The product of the operation is a result set with k elements. If k equals 1 (only one element resulted from applying
15 op to the response list), the multiplexor 22 sends this result as response Rp* to the client program 20. If k is greater than 1, an external mechanism, such as the graphical user interface discussed below, is employed to allow C to illustrate the list of results Rpl*...Rpk*.

20 Figures 4-6 illustrates one embodiment of the invention which provides temperature information from various sources (target instances of the server program). The example shown in Figures 4-6 is simplified to aid in the understanding of the invention. As would be known by one ordinarily skilled in

the art given this disclosure, the invention is not limited to temperature information systems, but instead is applicable to all types of computerized network systems which record, maintain and process information for any purpose. More specifically, in the example shown in Figures 4-6, each server program

5 includes information from one or more temperature sensors located in different locations/fixtures in a site, for example a store or warehouse. Each target server could represent a different floor or heating/cooling zone of a large building. Each zone could have multiple temperature sensors and multiple set points. Each of the

10 server programs could utilize different instances of the same program and share the same protocol. However, in this example is assumed that the client program and server program are only capable of communicating on a one-to-one basis. Conceptually, this example allows one client program to process a one-to-one request through the inventive multiplexor to each of the different instances of the

15 server program to return and control temperature information corresponding to multiple locations.

In this example, when a new value for a set point is selected by the user (e.g., a request by the client program), the temperature slowly changes from its current value until it reaches the desired set point. To read the current temperature, the user can click on a button to query the server (another request).

20 This follows the situation described above in which only a one-to-one protocol is available and a one-to-many capability desired.

Figures 4-6, are illustrated using a commercially available Web Browser (e.g., Netscape Navigator available from Netscape Communications, Inc., Palo Alto, California, USA). However, any such graphical user interface could be used.

5 In this example, the multiplexor behaves as a Web server, thus allowing users to connect to it with a Web browser. Figure 4 illustrates a list of potential instances of the server program (e.g., sneezy, snowwhite, cyberhealth, sleepy, grumpy, etc.) as item 2. The user would create the fan out list by checking one or more of the instances of the server program 2. In addition, item 3 in Figure 4
10 illustrates the check fields by which the user may select the operation op (e.g., list, aggregate, add, subtract, average, etc.) which is to be performed on the responses returned from the instances of the server program. Item 1 in Figure 4 is an entry field for the name of the service (e.g., a specific functionality provided by all server instances that can be queried by sending a request R_q) that is to be
15 multiplexed. In other words, the user would identify the server program in item number 1 of Figure 4.

Figure 5 illustrates a more customized page which displays current temperature and set points for two zones. While, in Figure 4, the server program could be selected by completing the input box 1, in Figure 5 the server program is
20 not a variable and is set as a temperature management program.

On the page shown in Figure 5, by typing a new value and clicking on the "Set" button 51, a new set point can be selected for either of two different items

(e.g., kitchen room temperature and refrigerator temperature). In this example it is assumed that the server programs for the kitchen and refrigerator are different instances of the same server program. To view how the change in the set point affects the current temperature, the user can click on "get current" and the page is repainted with the latest value 52.

This user interface provides a way to select the server application to be multiplexed 1, as well as the targets 2, 50 and operations 3. The multiplexor serves static pages (i.e. pages that are common to all instances of the server program for the same relative universal resource locator (URL)) from a local repository and manages requests for dynamic pages by queuing them and forwarding these requests to the servers according to some policy, e.g. immediately if all servers are online, as they become available, or at a scheduled time.

A request, generically called an "action", is saved if at least one of the targets 50 is not available. As shown in Figure 6, the invention can present the user with a list of similar actions 60 previously queued (or just the current one if it is the only one), including information about targets 61, the point in time when it was queued 62, and the current status (pending, ready, sent, etc.) 63. The user is given the option of either removing the action or queuing it to be processed as the targets become available. The user may select from the available results, and view results 64.

With the invention, for instance, the user could modify the set point 51 for this target (e.g., sleepy) 50 and click the "Set" button, thus sending the new request directly to that target. In fact, since the server "sleepy" is still online, it would receive the request immediately, and the display would be updated with the real time result.

To the contrary, the conventional systems would require modification of one or both target server programs (e.g., the kitchen temperature monitoring program and/or the refrigerator monitoring program) in order to allow such communication. In this example at least two servers are selected from the fan-out list. As shown in element 64 in Figure 6 there are two results, which indicates that at least two targets had been selected. The names of these servers would appear in the drop down list labeled "Targets". All servers (sleepy, grumpy, etc.) have information regarding kitchen and refrigerator, but each one corresponds to a different store or location.

While the overall methodology of the invention is described above, the invention can be embodied in any number of different types of systems and executed in any number of different ways, as would be known by one ordinarily skilled in the art, given this disclosure. For example, as illustrated in Figure 7, a typical hardware configuration of an information handling/computer system in accordance with the invention preferably has at least one processor or central processing unit (CPU) 700. For example, the central processing unit 700 could include various request processing units, mapping units, weighting units, adders,

subtractors, comparators, etc. Alternatively, as would be known by one ordinarily skilled in the art given this disclosure, multiple specialized CPU's (or other similar individual functional units) could perform the same processing, mapping, weighting, adding, subtracting, comparing, etc.

5 The CPU 700 is interconnected via a system bus 701 to a random access memory (RAM) 702, read-only memory (ROM) 703, input/output (I/O) adapter 704 (for connecting peripheral devices such as disk units 705 and tape drives 706 to the bus 701), communication adapter 707 (for connecting an information handling system to a data processing network) user interface adapter 708 (for
10 connecting a peripherals 709, 710 such as a keyboard, mouse, microphone speaker and/or other user interface device to the bus 701), a printer 712, and display adapter 713 (for connecting the bus 701 to a display device 714). The invention could be implemented using the structure shown in Figure 6 by including the inventive method within a computer program stored on the storage device 705.
15 Such a computer program would act on information supplied through the interface units 709, 710 or through the network connection 707. The system would then automatically produce the final desired product on the display 714, through the printer 712 or back to the network 707.

 The degree of transparency that can be achieved with the invention
20 depends greatly on the specific technology used to implement client program C, including the choice of protocol P. For example, a client program C having a graphical user interface implemented using object-oriented concepts known as

component model (such as systems implemented using the Java programming language and the JavaBeans specification, available from Sun Microsystems, Palo Alto, California, USA) may generate a component that can be included within other components. In a preferred embodiment, JavaBeans is used to encapsulate client C within a generic container that communicates with the multiplexor to iterate through the result set. Note that the invention can be advantageously used as a tool for the design of systems with certain characteristics.

The inventive multiplexor is similar to a proxy in that it appears to C as an instance of S, and at the same time appears to S as an instance of C. However, a major difference is that M establishes multiple connections for each request, whereas a proxy only establishes a connection with a single target.

Additionally, the inventive multiplexor may also encapsulate the complexity of communicating over sporadic connections (i.e., in systems in which not all S_i may be online at the time C sends a request). The invention is strictly independent of the programs C, S and the protocol P. By providing a building block with the buffer functionality, the development of applications that involve sporadic connections is greatly simplified with the invention.

As mentioned above, in the situation that involves sporadic connections, the target servers may not (all) be online at the time when the client attempts to send a request. The inventive multiplexor reduces the complexity of supporting this type of connection by storing requests in the buffer, dispatching requests as the target servers come online and storing the responses in the buffer until all

targets have been contacted. The operation op can then be applied and the resulting data (set) presented to the client program 20.

5 The invention uses a different approach than that used conventionally, which enables the use of the existing programs C and S, as well as protocol P, without modification, while accomplishing the goal of communicating with several instance of program Si. The invention can be exploited in different development situations, (e.g., to enable existing applications, or to provide a modular design for new applications).

10 If an application exists that cannot be modified in a cost-effective manner (e.g., because the source code or development skills used in the original application are no longer available, because of deployment and software distribution issues, or because of the need for backward compatibility), under certain conditions the inventive multiplexor can be used to transparently enhance the functionality so that multiple targets can be addressed simultaneously.

15 In the case in which a new application is being developed, the invention can be used as a building block for modular design. Thus, with the invention, several applications (multiple servers) can share a client program. Some standard protocols can then be commonly used to develop these applications. By providing support for these protocols, the inventive multiplexor can be used as a building
20 block.

The multiplexor is geared at communications between one client program C and multiple instances of the same server program S running on multiple

execution domains (e.g. nodes on a network). In general, programs C and S are built to work together by supporting a number of shared message types. A common design has servers maintaining state or store data, and clients retrieve this data on demand by sending a request Rq, usually in response to user actions.

5 While the same server program S can be instantiated (run) on several execution domains, the data that each instance of the server program S_i holds may be different and the response Rpi that each of the instances of the server program S_i sends to the client program C may thus be different.

The invention adapts existing, one-to-one client/server programs to handle
10 one-to-many communications transparently, without a need to modify the original software by modifying the request. This is extremely important as conventional applications are pushed beyond the limits of their originally intended use. The invention can be used as a building block in the development of new systems. By offloading the complexity of managing multicast communications, the invention
15 allows developers to concentrate on the specific requirements of the system being developed, resulting in significant cost and time savings.

By concentrating the functionality in one module, maintenance of the system (including debugging, extensions, etc.) can achieve economies of scale, since an enhancement can be leveraged by many systems.

20 The power of the inventive multiplexor is that it provides an almost transparent way to take a one-to-one protocol between a client and a server and extend it to multiple servers. For applications built from scratch, it reduces the

development effort by enabling the programmers to concentrate on how to handle the simpler case of one client-one server. In the case of legacy applications (defined not in terms of the technology they use, but rather as those that are very costly to modify) that only handle one-to-one communication, the use of a
5 multiplexor makes it possible to manage multiple connections in a manner transparent to the application.

The invention is especially applicable to networked embedded systems, in which a very large number of essentially identical devices interact over computer networks in complex ways. Internet-based applications such as electronic
10 commerce may also benefit from the invention by giving users the possibility of richer interactions without significant development effort. Traditional client/server applications can be re-deployed in a distributed fashion with the invention. The invention can also be used in the design of systems that do not involve network communications and is equally applicable to systems that require
15 processing asynchronous, identical function calls and that combine results.

While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.